

Automatic Vehicle Tracking and Recognition from Aerial Image Sequences

Ognjen Arandjelović
 School of Computer Science
 University of St Andrews
 St Andrews
 Fife KY16 9SX
 Scotland

E-mail: ognjen.arandjelovic@gmail.com

Abstract

This paper addresses the problem of automated vehicle tracking and recognition from aerial image sequences. Motivated by its successes in the existing literature we focus on the use of linear appearance subspaces to describe multi-view object appearance and highlight the challenges involved in their application as a part of a practical system. A working solution which includes steps for data extraction and normalization is described. In experiments on real-world data the proposed methodology achieved promising results with a high correct recognition rate and few, meaningful errors (type II errors whereby genuinely similar targets are sometimes being confused with one another). Directions for future research and possible improvements of the proposed method are discussed.

1. Introduction

In this paper we address the problem of automatic recognition of ground objects (mainly vehicles) from image sequences acquired from unmanned aircraft. This is a very challenging recognition scenario in which viewpoint and illumination can vary greatly, occlusions and background clutter are common, and data is of low quality. Generally speaking, target recognition systems comprise two distinct tasks, that of (i) representing object appearance and of subsequent (ii) representation matching. Both of these tasks are pervasive across different object recognition and matching problems, and have attracted a significant amount of research attention in the computer vision community [1, 2, 17]. This interest has particularly intensified in recent years, after significant advances towards practically viable systems have been made [7, 14, 17, 34].

The most prominent group approaches are local descriptor-based. Methods of this group employ descrip-

tors in a sparse fashion by focusing on a set of automatically localized interest points [19, 25]. When the number of detected interest points is large this approach can achieve impressive robustness to partial occlusion and pose [28]. However, a serious limitation of this approach is that it cannot deal well with untextured objects (sometimes referred to as smooth objects) which by their very nature do not exhibit appearance which results in a larger number of consistently well-localized interest points. This limitation has recently attracted increased research attention; shape-based approaches using boundary appearance features [14] or pure shape [6, 7] have demonstrated promising results on databases of objects which have distinct shapes. Other approaches include part-based methods, suitable to the recognition of articulated objects with distinctly recognizable parts (such as a face which can be seen as comprising two ‘eye parts’, a ‘mouth part’ etc. which vary in mutual configuration depending on the person’s pose and facial expression) [22]. A summary of different representation and matching techniques dominant in the existing literature is shown in Tables 1 and 2.

There are two key conceptual contributions we make to the current state-of-the-art. Firstly most existing methods address the problem at hand using individual images. In contrast in this paper the focus is on recognition from image sequences. In other words a sequence of frames/images of an unknown, query object is matched against a database of sequences of known, database targets. This problem setting is of an increasing significance considering the ease with which in our application image sequences can be acquired and stored. Secondly we show how a robust and pose-invariant system can be built by enriching the set of directly acquired exemplars with synthetically generated data, and how the resulting sets can be appropriately described and matched. We first summarize the key ideas and then explain each element of our method in detail in the next section.

Table 1. Classification of the most influential object recognition representations.

Object representations	Global representations	Appearance prototypes [27]
		Model-based [30]
	Local representations	Pictorial structure [22]
		Dense local features [25]
		Sparse local features [28]

Table 2. Classification of the most influential representation matching approaches.

Model matching	Single-view	Euclidean distance [27]
		Cosine distance [20]
	Multi-view	Single-view aggregation-based [23]
		Probability density-based [32]
		Embedded manifold-based [33]

Approach summary Considering the small scale of the target, model or part-based approaches are usually unsuitable for the problem in question. Notwithstanding its limitations, in this paper we focus on the use of a multi-view model in the form of a set of raw holistic appearance patches. This model motivates increasingly popular manifold-based approaches for matching, which exploit the structure of object’s appearance change as viewing parameters are varied. Following its success in related recognition problems and the potential for real-time performance in an incremental learning framework due to its computational efficiency and low storage requirements, here we specifically examine the use of canonical correlation analysis [11].

2. Method details

In this section we explain each of the constituent elements of the proposed algorithm. We start by summarizing the motivation and theory underlying the approach we adopt for matching sets of exemplars. Then we explain how exemplar data is extracted from raw imagery: Section 2.2.1 describes how a detected target is tracked through a sequence of images while possibly undergoing pose change, while Section 2.2.2 explains how viewpoint invariance is achieved by enriching the explicitly extracted set of exemplars with additional, synthetically generated images.

2.1. Overview of the baseline approach

For classification, canonical correlation analysis [26] is usually employed by computing canonical correlations (i.e. the cosines of principal angles) between linear manifolds [24, 31, 35]. Canonical correlations, $0 \leq \theta_1 \leq \dots \leq \theta_d \leq (\pi/2)$ between any two d -dimensional linear manifolds or hyperplanes \mathcal{L}_1 and \mathcal{L}_2 , are uniquely defined as [15]:

$$\cos \theta_i = \max_{\mathbf{u}_i \in \mathcal{L}_1} \max_{\mathbf{v}_i \in \mathcal{L}_2} \mathbf{u}_i^T \mathbf{v}_i \quad (1)$$

subject to $\mathbf{u}_i^T \mathbf{u}_i = \mathbf{v}_i^T \mathbf{v}_i = 1$, $\mathbf{u}_i^T \mathbf{u}_j = \mathbf{v}_i^T \mathbf{v}_j = 0$, $i \neq j$. The solution can be obtained by applying the Singular Value Decomposition (SVD) [18], whose complexity is $O(d^3)$ where d is the dimensionality of the manifolds (d is typically small). Each image set is represented by a linear manifold and the angles between two low-dimensional manifolds are exploited as a similarity measure between two image sets. The canonical vectors in each pair are visually similar despite the large changes of object pose. The first pair of canonical vectors corresponds to the most correlated vectors, each of which is spanned by any linear combination of the respective image set. The next pairs of canonical vectors represent the directions of the next most similar data variations of the two sets in other dimensions. CCA effectively finds common modes (e.g. target object pose) of two image sets.

2.2. Data extraction

In designing a practical object recognition system it is of crucial importance to appreciate that the actual recognition algorithm operates in the context of the preceding data acquisition and extraction stages. Specifically, in the task of target recognition from aerial images acquired using unmanned aircraft, the target needs to be located and tracked. On the lowest level, this is done by the on-board control system which employs the GPS coordinates of the target and the plane, and a gyroscope. Difficulties are already introduced here – the target moves widely across the image, as shown in Figure 1, often even entirely disappearing from the view. Given that vehicles appear in aerial images mostly as smooth object, resulting in few stably detectable keypoints, we found local-feature-based tracking [29] unreliable. Instead we propose an alternative solution which comprises two steps. Firstly, the problem of initializing tracking by detection is solved by registering consecutive image frames globally. This is readily achieved because video sequences of interest in this paper are mostly comprised of static backgrounds. Hence in the context of registration of frames the target can be considered an outlier (in that it moves from a frame to frame). After consecutive frames are registered, the target itself is easily detected by performing simple background subtraction. To account for noise, morphological image processing is performed to

remove spurious regions of frame-to-frame difference, with the correct region reliably being detected as the one with the largest contiguous area over which the aforementioned difference exceeds a threshold.



Figure 1. Raw input frame showing a poorly localized target by the aircraft’s on-board system.

2.2.1 Target tracking

Following the localization of the target, we track it until it is no longer entirely visible. For this we employ the well known Lucas-Kanade tracker [3, 16], with 6 affine degrees of freedom. In this algorithm, at each image-to-image transition in a sequence, the generalized position vector is iteratively updated by minimization of the error (i.e. appearance difference in Euclidean distance sense) between the region of an image at that position and the warped template of the target from the preceding image in the sequence. In more detail, given the appearance of the target $I_i(x, y)$ in the i -th image in a sequence, the corresponding region of interest in the subsequent frame I_{i+1} is found by minimizing:

$$\sum_{(x,y) \in \mathcal{R}} [I_{i+1}(\mathcal{W}(x, y; \mathbf{p})) - I_i(x, y)], \quad (2)$$

where \mathcal{R} is the quadrilateral region of interest specifying the target in I_i , \mathcal{W} the warping function (an affine warp in our case), and \mathbf{p} the warp parameters. The optimal \mathbf{p} is found through iterative descent by linearizing $I_{i+1}(\mathcal{W}(x, y; \mathbf{p}))$ using Taylor expansion, giving:

$$\sum_{(x,y) \in \mathcal{R}} \left[I_{i+1}(\mathcal{W}(x, y; \mathbf{p}_j)) + \nabla I \frac{\partial \mathcal{W}}{\partial \mathbf{p}} \Delta \mathbf{p} - I_i(x, y) \right], \quad (3)$$

where \mathbf{p}_0 are the initial warp parameters, and $\mathbf{p}_1 \dots \mathbf{p}_j$ the iterative refinement sequence. We found this approach to

work well on our data set and expect a similar level of performance on imagery acquired under similar conditions (elevation and angle to the target).

2.2.2 Appearance set generation

Ideally, the target recognition system receives views of the target across 360° range, obtained by a circular reconnaissance manoeuvre over the target. However, as a consequence of the difficulties involved in the camera’s control system locking onto the target (see Section 2.2), the range of views obtained from each flyover is incomplete. Furthermore, in the experiments reported in this paper, it was decided to use each tracking burst (from the initial detection of the target until the target is lost as described in Section 2.2.1) as a single training set, thus restricting the range of views available for training. The reason for this lies in the problem of concatenating aforementioned tracking bursts into one data set without introducing training data artefacts. These include, for example, repeated views of the target or differing aeroplane viewpoint angle due to multiple flyovers.

The described fragmentation of input video sequences is a serious problem as the baseline method of interest offers no invariance in this regard. As described in Section 2.1, for canonical correlations to extract meaningful similarity between two data sets, a common form of variation must exist. While robust to the presence of dissimilar data, whether due to different viewing conditions or noise, true rotational or view invariance is not inherent in this approach.

Instead of attempting to achieve invariant matching, in this paper we examine an alternative approach whereby invariance is achieved explicitly, by synthetic data augmentation. We summarize the key stages in the proposed method:

1. **Re-warping:** Our tracking approach involves the estimation of pose parameters for the target. To quasi-normalize this view, we “re-warp” the target to fit its initial pose (i.e. the pose in the first frame as explained in Section 2.2), as shown in Figure 2.
2. **Full view generation:** The “re-warped” image of the target is now synthetically rotated across 360° at 10° intervals, as shown in Figure 3.
3. **Subspace estimation:** Synthetic views from all tracked templates are compiled and used to estimate a linear subspace, which is used as the final representation of target’s appearance in the video.

3. Empirical evaluation

In the empirical evaluation reported in this paper, we used data acquired during five flights and multiple target




Figure 2. The first step of our data processing and normalization involves “re-wrapping” each tracked patch to fit the initial pose.




Figure 3. Seven synthetically generated views of the target, produced by rotating the “re-wrapped” patch (see Figure 2) across different angles.

Table 3. A summary of data obtained during the five flights which were used for empirical evaluation.

Flight	Date	Images	Resolution
1	28 October 2009	28364	1360×1024
2	06 November 2009	34262	1360×1024
3	10 November 2009	17636	1360×1024
4	12 November 2009	14076	1360×1024
5	17 November 2009	21108	1360×1024

flyovers. The details of each session are summarized in Table 3.

To test the proposed recognition system we used six different targets extracted from these five flights. Three views of each and their symbolic names are shown in Figure 4. The scale of each target varied throughout input video, depending on its geometry (significantly different for a car and a tent, for example), camera viewpoint, and aeroplane dis-

tance (mainly affected by its altitude). The size of the target image patch diagonal lied in the range of 70 to 200 pixels. We normalize tracked target patches by warping them to the uniform scale of 100×100 pixels.




Figure 4. Six different targets, shown in three different poses each, used to test the proposed algorithm.

3.1. Results

Target recognition was performed by matching a novel, query data set against training sets of each of the six targets. It was identified as the target that it matched with the highest confidence. The results of rank-1 recognition are summarized using the confusion matrix in Table 4. As the confusion matrix illustrates, the proposed method correctly identified the novel target in all but a few cases. An

Table 4. The target confusion matrix obtained in our experiments (shown is error rate in %). Our algorithm made few errors (wrong target assignments) and the few that were made can be seen to correspond to genuinely similar objects.

	Car 1	Tent	Van 1	Car 2	Van 2	Car 3	Total
Car 1	—	0.0	3.4	0.0	0.0	4.7	7.0
Tent	0.0	—	0.0	0.0	0.0	0.0	0.0
Van 1	2.2	0.0	—	0.0	0.0	1.5	3.7
Car 2	0.0	0.0	0.0	—	1.9	0.0	1.9
Van 2	0.0	0.0	0.0	2.9	—	0.0	2.9
Car 3	6.7	0.0	0.8	0.0	0.0	—	7.5

inspection of incorrect target assignments readily shows a clear structure of such errors, targets with genuinely similar appearance being occasionally confused with one another (e.g. car 2 and van 2 which are both dark, of similar shape, and only a small difference in size which is not readily apparent from images considering that the camera-target distances are unknown).

We further tested the sensitivity of our method to the amount of training data. As explained in Section 2.2.2, in principle even a single image of the target can be used for matching, as a synthetic, compatible 360° set is generated from each image. Thus, we explored how the correct recognition rate is affected by gradual removal of an increasing amount of tracked templates. Figure 5 summarizes the results obtained and shows that the proposed method exhibits slow, graceful performance degradation.




Figure 5. Decay of correct identification rate as the amount of data used for training is reduced. Graceful degradation is demonstrated with a high recognition rate even when 40% of the data is discarded.

4. Conclusions

This paper presented preliminary experiments and findings for target recognition from unmanned reconnaissance aircraft using a method based on canonical correlations, a well known statistical method for comparing sets of high

dimensional vectors. A framework for employing canonical correlations in this scenario was described, followed by a description of experiments conducted to assess its effectiveness. These preliminary results show promising performance with high correct recognition rate and graceful performance decay in the presence of a reducing amount of data.

4.1. Future work

Results reported here encourage and call for more experimental and research effort in the development of a canonical correlations-based target recognition system. These include:

- **Data variation modelling:** In this paper we only evaluated the performance of canonical correlations using linear subspaces. This approach is almost always inferior to one which takes into account the nonlinear nature of object appearance manifolds. Specifically, we would like to investigate the performance of a method which would represent each 360° appearance variation as a set of subspaces, which are then mutually compared in a manner similar to that described in [10]. Another promising direction involves the use of probabilistic extensions of canonical correlations [9].
- **Different appearance representation:** Here we only explored the use of raw image appearance to model target appearance. The use of more complex representations, e.g. based on oriented gradients [28, 21, 5], robust edges [12, 13, 8] or colour invariants [4], could result in increased robustness of the method and possibly remove the need for synthetic data augmentation over full 360° .

References

- [1] G. Aggarwal and D. Roth. Learning a sparse representation for object detection. *In Proc. European Conference on Computer Vision*, 2002.
- [2] A. Ahmadyfard and J. Kittler. A comparative study of two object recognition methods. *In Proc. British Machine Vision Conference*, pages 1–10, 2002.
- [3] O. Arandjelović. Contextually learnt detection of unusual motion-based behaviour in crowded public spaces. *In Proc. International Symposium on Computer and Information Sciences*, pages 403–410, 2011.
- [4] O. Arandjelović. Colour invariants under a non-linear photometric camera model and their application to face recognition from video. *Pattern Recognition*, 45(7):2499–2509, 2012.
- [5] O. Arandjelović. Gradient edge map features for frontal face recognition under extreme illumination

- changes. In *Proc. British Machine Vision Conference*, 2012. DOI: 10.5244/C.26.12.
- [6] O. Arandjelović. Matching objects across the textured-smooth continuum. In *Proc. Australasian Conference on Robotics and Automation*, pages 354–361, 2012.
- [7] O. Arandjelović. Object matching using boundary descriptors. In *Proc. British Machine Vision Conference*, 2012. DOI: 10.5244/C.26.85.
- [8] O. Arandjelović. Making the most of the self-quotient image in face recognition. In *Proc. IEEE International Conference on Automatic Face and Gesture Recognition*, 2013. DOI: 10.1109/FG.2013.6553708.
- [9] O. Arandjelović. Discriminative extended canonical correlation analysis for pattern set matching. *Machine Learning*, 94(3):353–370, 2014.
- [10] O. Arandjelović and R. Cipolla. Face set classification using maximally probable mutual modes. In *Proc. IAPR International Conference on Pattern Recognition*, pages 511–514, 2006.
- [11] O. Arandjelović and R. Cipolla. Achieving robust face recognition from video by combining a weak photometric model and a learnt generic face invariant. *Pattern Recognition*, 46(1):9–23, 2013.
- [12] O. Arandjelović, D. Pham, and S. Venkatesh. Efficient and accurate set-based registration of time-separated aerial images. *Pattern Recognition*, 2015. DOI: 10.1016/j.patcog.2015.04.011.
- [13] O. Arandjelović, D. Pham, and S. Venkatesh. Group-wise registration of aerial images. In *Proc. International Joint Conference on Artificial Intelligence*, 2015.
- [14] R. Arandjelović and A. Zisserman. Smooth object retrieval using a bag of boundaries. In *Proc. IEEE International Conference on Computer Vision*, pages 375–382, 2011.
- [15] F. Bach and M. I. Jordan. Kernel independent component analysis. *Journal of Machine Learning Research*, 3:1–48, 2002.
- [16] S. Baker and I. Matthews. Lucas-kanade 20 years on: A unifying framework. *International Journal of Computer Vision*, 56(3):221–255, 2004.
- [17] P. N. Belhumeur and D. J. Kriegman. What is the set of images of an object under all possible illumination conditions? *International Journal of Computer Vision*, 28(3):245–260, 1998.
- [18] Å. Björck and G. H. Golub. Numerical methods for computing angles between linear subspaces. *Mathematics of Computation*, 27(123):579–594, 1973.
- [19] M. Bregonzio, T. Xiang, and S. Gong. Fusing appearance and distribution information of interest points for action recognition. *Pattern Recognition*, 45(3):1220–1234, 2012.
- [20] W. Chen, M. J. Er, and S. Wu. Illumination compensation and normalization for robust face recognition using discrete cosine transform in logarithm domain. *IEEE Transactions on Systems, Man, and Cybernetics – Part B: Cybernetics*, 36(2):458–466, 2006.
- [21] N. Dalai and B. Triggs. Histograms of oriented gradients for human detection. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 1:886–893, 2005.
- [22] P. F. Felzenszwalb and D. Huttenlocher. Pictorial structures for object recognition. *International Journal of Computer Vision*, 61(1):55–79, 2005.
- [23] V. Ferrari, T. Tuytelaars, , and L. Van Gool. Integrating multiple model views for object recognition. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2:105–112, 2004.
- [24] K. Fukui and O. Yamaguchi. Face recognition using multi-viewpoint patterns for robot vision. *International Symposium of Robotics Research*, pages 192–201, 2003.
- [25] M. Heikkilä, M. Pietikäinen, and C. Schmid. Description of interest regions with local binary patterns. *Pattern Recognition*, 42(3):425–436, 2009.
- [26] H. Hotelling. Relations between two sets of variates. *Biometrika*, 28:321–372, 1936.
- [27] Y. LeCun, F. J. Huang, and L. Bottou. Learning methods for generic object recognition with invariance to pose and lighting. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2:97–104, 2004.
- [28] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2003.
- [29] R. Martin and O. Arandjelović. Multiple-object tracking in cluttered and crowded public spaces. In *Proc. International Symposium on Visual Computing*, 3:89–98, 2010.
- [30] A. S. Mian, M. Bennamoun, and R. Owens. Three-dimensional model-based object recognition and segmentation in cluttered scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(10):1584–1601, 2006.
- [31] M. Nishiyama and O. Yamaguchi. Face recognition using the classified appearance-based quotient image. In *Proc. IEEE International Conference on Automatic Face and Gesture Recognition*, pages 49–54, 2006.

- [32] G. Shakhnarovich, J. W. Fisher, and T. Darrel. Face recognition from long-term observations. *In Proc. European Conference on Computer Vision*, 3:851–868, 2002.
- [33] A. Shashua, A. Levin, and S. Avidan. Manifold pursuit – a new approach to appearance based recognition. *In Proc. IAPR International Conference on Pattern Recognition*, pages 590–594, 2002.
- [34] J. Sivic, B. Russell, A. Efros, A. Zisserman, and W. Freeman. Discovering object categories in image collections. *In Proc. IEEE International Conference on Computer Vision*, pages 370–377, 2005.
- [35] L. Wolf and A. Shashua. Learning over sets using kernel principal angles. *Journal of Machine Learning Research*, 4(10):913–931, 2003.